

KELLEY DRYE & WARREN LLP

A LIMITED LIABILITY PARTNERSHIP

WASHINGTON HARBOUR, SUITE 400

3050 K STREET, NW

WASHINGTON, D.C. 20007-5108

(202) 342-8400

FACSIMILE

(202) 342-8451

www.kelleydrye.com

NEW YORK, NY
TYSONS CORNER, VA
CHICAGO, IL
STAMFORD, CT
PARSIPPANY, NJ

BRUSSELS, BELGIUM

AFFILIATE OFFICES
JAKARTA, INDONESIA
MUMBAI, INDIA

DIRECT LINE: (202) 342-8518

EMAIL: lcohen@kelleydrye.com

August 12, 2010

VIA ECFS

Marlene H. Dortch
Secretary
Federal Communications Commission
The Portals
445 - 12th Street, SW
Washington, DC 20554

Re: Notice of *Ex Parte* Presentation, *Preserving the Open Internet*, GN
Docket No. 09-191, and *Broadband Industry Practices*, WC Docket No.
07-52

Dear Ms. Dortch:

Yesterday, representatives of the Fiber-to-the-Home Council ("FTTH Council") met, in three separate meetings, with the following staff of the Commission:

Office of Engineering and Technology – Julius Knapp, Chief, and Ronald Repasi, Deputy Chief; Office of Strategic Planning and Policy Analysis – Douglas Sicker;

Office of Commissioner Michael Copps – Jennifer Schneider, Senior Policy Advisor and Legal Advisor for Broadband, Wireline and Universal Service;

Office of Commissioner Mignon Clyburn – Angela Kronenberg, Acting Chief of Staff and Wireline Legal Advisor.

The FTTH Council representatives were: Joseph Savage (President, FTTH Council), Mike Hill (Vice Chairman, FTTH Council), Kevin Morgan (Adtran), Michael Johnston (Jackson Energy Authority), Dale Merten (The Toledo Telephone Company), Bryan Geiger (GVTC Communications), Ashley Phillips (EATEL), and myself.

Marlene H. Dortch
August 12, 2010
Page Two

The purpose of the meeting was to discuss the attached presentation, which reviewed *A Network Engineer's Primer on Broadband Internet Access Services and Reasonable Network Management Practices for Wireline Networks*. This primer was submitted by the FTTH Council on January 14, 2010 as an attachment to its comments. The following summarizes the "bright-line" rules proposed by the FTTH Council to govern determinations about what constitutes reasonable network management practices for wireline networks should the Commission proceed to adopt its proposed rules:

- **Should the government decide to adopt any rules concerning network management, they should apply only to the provision of Broadband Internet Access Service and not to the offering, operation, or management of any other services, including managed or specialized services, offered by the provider even if such services are provided over common facilities.** Modern communications networks are used to provide a wide (and constantly changing) array of services, each of which is comprised of features and functionality sought by users and each of which accordingly places different demands on the network. In determining whether and how to provide these services, network operators weigh numerous factors and make complex calculations about whether to invest substantial amounts of capital to deploy additional capacity or manage within current capacity. Further, a balance made by one entity likely is irrelevant (or at best of minimal relevance) to a balance made by another – after all, today's communications networks vary tremendously, even among providers offering the same or similar arrays of services.
- **Broadband Internet access providers may offer users different classes of services with varying performance capabilities, which may include the imposition of limitations on users' bandwidth and consumption subject to full and understandable disclosure.** It is a common and well-accepted practice today for broadband providers to offer classes (tiers) of service, which are sold at different prices for different performance capabilities (e.g. bandwidth) and with different terms and conditions (e.g. limitations on consumption). The key is to make sure users are given sufficiently clear information about these classes of service – both capabilities and limitations – so they can make informed decisions.
- **Broadband Internet access providers may provide quality of service or other performance guarantees to providers and users of applications so long as they are not offered on an exclusive basis for any particular application category.** Many Internet applications are highly-sensitive to latency and jitter and packet loss – and newer real-time video applications also must have sufficient bandwidth to provide a satisfactory experience for users. Providers of applications and users of these applications would benefit from Quality of Service ("QoS") guarantees and, in fact, often request them. (Many of these providers already use Content Delivery Networks ("CDNs") to ensure QoS from the host server to a server collocated with the broadband Internet access provider.) So long as any QoS guarantees are offered on a

Marlene H. Dortch

August 12, 2010

Page Three

non-exclusive basis – that is to all providers of any particular category of application – they should be permitted.

- **Broadband Internet access providers may choose to increase or alter network capacity by building additional facilities or installing or tuning electronics -- and the government can provide incentives to build additional capacity¹ -- but providers should not be required to do so.** Decisions regarding increasing network capacity are exceedingly complex, involving not just network engineering concerns but all the factors that are part of a business case. In addition, these decisions often have to be made relatively rapidly given the growth of Internet traffic. Further, the circumstances that lead to the possible investment in more capacity may be driven by a small subset of subscribers. Finally, it would be unprecedented for the government to mandate or prevent the construction of facilities and equipment by communications providers other than those who have a monopoly.

- **Broadband Internet access providers may throttle and shape Internet access traffic (1) so long as they do so for all applications of a particular type and so long as the action is of a short duration and not repeatedly imposed; (2) if it is imposed on a individual user, so long as such action is part of the terms and conditions agreed upon by the user when choosing the service, or (3) in emergency situations .** Because of the potential high volatility of Internet traffic, network managers often face surges beyond those anticipated by and engineered into the network. To ensure a satisfactory level of service, network managers should be able to address these events in several ways. First, they may throttle or shape traffic if it applies to all applications of a particular type. Second, they may throttle or shape traffic for a particular user if the provider informs the user in clear and understandable terms in advance that such activity may occur. Third, they may throttle or shape traffic if emergency situations – to permit critical traffic to get through or to prevent the network from crashing.

- **Broadband Internet access providers may permit CDNs or others, including themselves or affiliated entities, to install and have traffic directed towards local caching equipment.** CDNs are so widespread and have such generally recognized benefits for users and content and applications providers that they should be encouraged – even though they may give advantages to particular entities (as opposed to particular applications).

- **Broadband Internet access providers may block or otherwise restrict spam, malware, and similar traffic harmful to the network and unlawful traffic (or unlawful transfer of content) – as well as provide priorities for emergency traffic and secure**

¹ The FTTH Council, in fact, believes the government should be involved in providing incentives to accelerate deployment of FTTH and other next-generation access networks, and it has consistently advocated for measures to accomplish this objective, including investment tax credits and tax credit bonds.

KELLEY DRYE & WARREN LLP

Marlene H. Dortch
August 12, 2010
Page Four

transmissions. There is a virtually unanimous agreement that these activities are either so harmful in the case of spam and malware or beneficial in the case of emergency and secure traffic that network managers should be able to address them expeditiously by use of practices generally accepted in the industry.

- **Broadband Internet access providers should publicly disclose information about network usage limitations or conditions and general information about their network management practices but may limit disclosure of network management practices that would negate or otherwise undermine any permissible network management practice.** Providing users and applications and content providers with sufficient and understandable information about network usage limitations or conditions is important to ensure they make informed choices about subscribing to the service; however, providers should not be ordered to disclose information regarding the precise nature of any network management practices, including technical methods employed to manage traffic. Such granular detail about network management practices is often proprietary, crucial to preserve network security and reliability, and is constantly changing in response to new congestion problems or malware attacks.

Should you have questions about these meetings, please contact me.

Sincerely,



Thomas Cohen
Kelley Drye & Warren LLP
3050 K Street, NW
Suite 400
Washington, DC 20007
Tel. (202) 342-8518
Fax. (202) 342-8451
tcohen@kelleydrye.com
Counsel for the Fiber to the Home Council

Attachment: (1) *Fiber-to-the-Home Council Presentation on Reasonable Network Management Practices for Wireline Networks, August 11, 2010, GN Docket 09-191 and WC Docket 07-52*
(2) *A Network Engineer's Primer on Broadband Internet Access Services and Reasonable Network Management Practices for Wireline Networks, January 14, 2010.*

KELLEY DRYE & WARREN LLP

Marlene H. Dortch
August 12, 2010
Page Five




cc: Julius Knapp
Ronald Repasi
Douglas Sicker
Jennifer Schneider
Angela Kronenberg

Fiber-to-the-Home Council

Presentation on Reasonable Network Management Practices for Wireline Networks

August 11, 2010
(GN Dockets 09-191 and WC Docket 07-52)

TAP INTO THE MOST VALUABLE BROADBAND RESOURCE AVAILABLE



Building Fiber-to-the-Home
Communities Together



Introduction: FTTH Council Network Engineers' Primer

- Filed by the FTTH Council on January 14, 2010
- Endorsing Engineers:
 - Michael Johnston, Vice President of IT and Broadband, Jackson Energy Authority
 - Dale Merten, Chief Operating Officer, The Toledo Telephone Company
 - George O'Neal, Vice President Network Services, (Bryan Geiger, NOC Manager) GVTC Communications
 - Dan Pecarina, Vice President of Technology, Hiawatha Broadband Communications
 - Ashley Phillips, EATEL Director of Network Engineering & Operations
 - Stephen Schneider, Director of Engineering, (Rod Kutemeier, General Manager) Sunflower Broadband
- Objective: Provide a set of "bright line" rules the Commission should employ should it decide to intervene in network management activities for Broadband Internet Access Services

Modern Wireline Networks: Complex and Dynamic

- **Modern Wireline Networks Provide a Wide Array of Services, of which Broadband Internet Access is Only One --**
 - Voice – Circuit & IP
 - Broadband Internet Access Services
 - Video – RF or IPTV
 - Carrier-grade services for business customers
 - Wireless backhaul for mobile providers
 - Wholesale “lit” capacity to other carriers

Modern Wireline Networks: Complex and Dynamic

- **With Each Service Having Different Network Requirements to Ensure Users have a Quality Experience –**
 - Voice often requires very low latency, jitter, and packet-loss
 - Many services dependent on Internet access, *e.g.* email and web searches, can tolerate more jitter and latency; Others, *e.g.* streaming video, require substantial bandwidth, low latency and jitter and very low packet-loss
 - Carrier-grade services usually have SLAs
 - Mobile backhaul requires extremely low latency and jitter and timing synchronization

Modern Wireline Networks: Complex and Dynamic

- **Thus Providers need to Engage in Inter-Service and Intra-Service Network Management --**
 - Inter-Service Examples
 - TDM Voice given priority over broadband Internet access traffic to avoid problems caused by bursts in Internet transmissions
 - IPTV requires allocated bandwidth to lessen interference from large bursts of Internet traffic
 - Business or public safety services demand allocated bandwidth to ensure higher QoS
 - Intra-Service Example for Broadband Internet Access Service
 - Over-the-top VoIP is given priority

Modern Wireline Networks: Complex and Dynamic

- **But Local Access Providers Only Control Part of the End-to-End Path and for Broadband Internet Access Service Transmission Quality is Affected at Many Points --**
 - Customer Premises, including computers and local area networks
 - Local Access Network (with at least some shared facilities)
 - Content Distribution Network and other Caching
 - Transport to the Internet (Middle-Mile)
 - Internet Backbone
 - Content/Applications Servers

Modern Wireline Networks: Broadband Internet Access Service

- **Characteristics of Broadband Internet Access Traffic –**
 - “Best Efforts” Network where the control of data flow is shared by end stations and network providers
 - End Station control facilitates development of new content/applications, *e.g.* P2P, Online Gaming, Cloud Computing
 - That has led to the enormous growth in traffic (doubling every 2 years)
 - But innovative applications can stress networks
 - And openness facilitates malicious traffic
 - And government imposes oversight requirements

Network Management Techniques Employed by Broadband Internet Access Providers

- **Increase Network Capacity**
 - Important tool but capital intensive
- **Establish Classes (Tiers) of Service for Users**
 - Requires transparency/disclosure, education and customer responsiveness
 - May still require use of other management techniques

Network Management Techniques Employed by Broadband Internet Access Providers

- **Direct Management of Traffic**

- Oversubscription of Shared Facilities
- Traffic Control Techniques (with Policy Management) –
 - General -- Throttling with no reference to underlying application or the source/destination to address a specific event; If cause of congestion is specific application, limit use or give other applications higher priority
 - DiffServe and MPLS
 - Malware Detection and Blocking
 - Deep Packet Inspection
 - Policy Control

- **Content Delivery Networks and Local Caching**

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

1. IF GOVERNMENT OVERSIGHT ORDERED, FOCUS SHOULD ONLY BE ON BROADBAND INTERNET ACCESS SERVICE

Should the government decide to adopt any rules concerning network management, they should apply only to the provision of Broadband Internet Access Service and not to the offering, operation, or management of any other services, including managed or specialized services, offered by the provider even if such services are provided over common facilities.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

2. PERMIT OFFERING OF DIFFERENT CLASSES OF SERVICE TO USERS

Broadband Internet access providers may offer users different classes of services with varying performance capabilities, which may include the imposition of limitations on users' bandwidth and consumption subject to full and understandable disclosure.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

3. PERMIT QoS SO LONG AS NON-EXCLUSIVE

Broadband Internet access providers may provide quality of service or other performance guarantees to providers and users of applications so long as they are not offered on an exclusive basis for any particular application category.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

4. DO NOT REQUIRE, BUT INSTEAD PROVIDE INCENTIVES FOR, INCREASES IN NETWORK CAPACITY OR EFFICIENCY

Broadband Internet access providers may choose to increase or alter network capacity by building additional facilities or installing or tuning electronics -- and the government can provide incentives to build additional capacity -- but providers should not be required to do so.

*The FTTH Council, in fact, believes the government should be involved in providing incentives to accelerate deployment of FTTH and other next-generation access networks, and it has consistently advocated for measures to accomplish this objective, including investment tax credits and tax credit bonds.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

5. TRAFFIC MAY BE THROTTLED/SHAPED UNDER LIMITED CIRCUMSTANCES

Broadband Internet access providers may throttle and shape Internet access traffic

(1) so long as they do so for all applications of a particular type and so long as the action is of a short duration and not repeatedly imposed,

(2) if it is imposed on a individual user, so long as such action is part of the terms and conditions agreed upon by the user when choosing the service, or

(3) in emergency situations.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

6. PERMIT LOCAL CACHING

Broadband Internet access providers may permit CDNs or others, including themselves or affiliated entities, to install and have traffic directed towards local caching equipment.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

7. PROVIDE WIDE LATITUDE TO DEAL WITH MALWARE AND ENSURE THE FLOW OF EMERGENCY AND SECURE TRAFFIC

Broadband Internet access providers may block or otherwise restrict spam, malware, and similar traffic harmful to the network and unlawful traffic (or unlawful transfer of content) – as well as provide priorities for emergency traffic and secure transmissions.

Reasonable Network Management Practices for Wireline Networks: Bright-Line Rules

8. REQUIRE REASONABLE DISCLOSURE OF NETWORK MANAGEMENT PRACTICES

Broadband Internet access providers should publicly disclose information about network usage limitations or conditions and general information about their network management practices but may limit disclosure of network management practices that would negate or otherwise undermine any permissible network management practice.

**A NETWORK ENGINEER'S PRIMER
ON BROADBAND INTERNET ACCESS SERVICES AND
REASONABLE NETWORK MANAGEMENT PRACTICES
FOR WIRELINE NETWORKS**

ENDORSING ENGINEERS:

Michael Johnston, Vice President of IT and Broadband, Jackson Energy Authority

Dale Merten, Chief Operating Officer, The Toledo Telephone Company

George O'Neal, Vice President Network Services, GVTC Communications

**Dan Pecarina, Vice President of Technology, Hiawatha Broadband
Communications**

Ashley Phillips, EATEL Director of Network Engineering & Operations

Stephen Schneider, Director of Engineering, Sunflower Broadband

**Filed with the Comments of the Fiber-to-the-Home Council
in Federal Communications Commission GN Docket 09-191 and WC Docket 07-52**

January 14, 2010

TABLE OF CONTENTS

	Page
I. EXECUTIVE SUMMARY	1
II. INTRODUCTION	5
III. BASICS OF NETWORK MANAGEMENT.....	7
A. Traditional Telephone Networks	7
B. Management of Modern Communications Networks.....	8
IV. NETWORK MANAGEMENT OF BROADBAND INTERNET ACCESS SERVICES.....	10
A. Overview of Internet Traffic.....	10
1. Introduction.....	10
a. Transmission Control Protocol	10
b. The Internet as a “Best Efforts” Network.....	11
2. Types of Traffic and Their Requirements.....	12
a. Beneficial Traffic	12
b. Malicious traffic.....	13
c. Legal Requirements	15
3. Traffic Growth and Evolution.....	15
B. Points and Sources of Congestion for Internet Traffic	16
C. Network Management Techniques Employed Today by Broadband Internet Access Providers	18
1. Introduction: Users and Providers Address Congestion	18
2. Increase Network Capacity	18
3. Establish Classes (Tiers) of Service for Users.....	20
4. Direct Management of Traffic	20
a. Oversubscription of Shared Facilities.....	20
b. Traffic Control Techniques.....	21
(1) General Traffic Control Techniques	21
(2) DiffServ and MPLS	22
(3) Malware Detection and Blocking	22
(4) Deep Packet Inspection.....	23
(5) Policy Control	23
c. Content Delivery Networks and Local Caching	25

TABLE OF CONTENTS
(continued)

Page

V.	REASONABLE NETWORK MANAGEMENT PRACTICES FOR WIRELINE NETWORKS	26
----	--	----

**A NETWORK ENGINEER'S PRIMER
ON BROADBAND INTERNET ACCESS SERVICES AND
REASONABLE NETWORK MANAGEMENT PRACTICES
FOR WIRELINE NETWORKS**

ENDORSING ENGINEERS:

Michael Johnston, Vice President of IT and Broadband, Jackson Energy Authority

Dale Merten, Chief Operating Officer, The Toledo Telephone Company

George O'Neal, Vice President Network Services, GVTC Communications

**Dan Pecarina, Vice President of Technology, Hiawatha Broadband
Communications**

Ashley Phillips, EATEL Director of Network Engineering & Operations

Stephen Schneider, Director of Engineering, Sunflower Broadband

**Filed with the Comments of the Fiber-to-the-Home Council
in Federal Communications Commission GN Docket 09-191 and WC Docket 07-52**

January 14, 2010

I. EXECUTIVE SUMMARY

Modern communications networks and operations are inherently complex:

- Today's networks simultaneously carry voice traffic using traditional TDM or Internet-based IP ("VoIP") protocols, data traffic via TDM technology or any number of packet offerings, including IP, some of which flow to and from the Internet, and video traffic in "cable" or "Radio Frequency" format, IPTV, or IP video streamed or sent as large files to and from the Internet.
- They handle widely varying usage profiles that place a broad range of capacity and latency demands.
- They are constantly threatened by malware and spam.
- They need to handle numerous legal requirements, including law enforcement assistance and copyright infringement enforcement.

In this environment, local network managers must address rapidly changing services or applications with accompanying differences in network flows and requirements, interconnections with and hand-offs to and from different networks, constantly mutating and more malicious malware, legal requirements to protect content, law enforcement mandates, and emergency communications. Further, while congestion may occur on local networks, it is quite possible that perceived congestion and

degradation in service may be determined by events and network choices of third parties that occur before content is even delivered to local networks. Moreover, end users themselves may experience congestion because of deficiencies in their own terminal equipment, inadequate computer processing power while running capacity-hungry or multiple applications, and bottlenecks within customer operated home or business networks.

Network managers, of course, can remedy, at least temporarily, some of the issues that arise on their networks by deploying additional capacity. That is the benefit of fiber-to-the-premises access networks, which have tremendous capacity, reduced dependence on shared facilities, and can be readily upgraded. But, given the growth and evolution of video communications and spam and malware, even local access networks based on fiber architectures need to address the trade-off between the creation of new capacity and load management. In addition, capacity increases will address only a subset of possible issues to efficient network management. Thus, traffic management remains essential.

The Federal Communications Commission now proposes the adoption of six network neutrality rules which the agency tentatively concludes is necessary to ensure users of broadband Internet access services can access content and applications and connect devices of their choice and to prohibit providers of these services from discriminating against any content, applications or services. The proposed rules would allow providers to engage in "unspecified" reasonable network management practices. Thus there is a need for clear guidance as to what types of network management practices are "reasonable" and will be allowed. In addition, because broadband Internet access services are just one of the many services offered simultaneously over modern communications networks and thus the traffic patterns associated with Internet access and use are only a subset of the elements engineers must factor into overall network management calculations, it is important that the Commission tread carefully in establishing any limitations on traffic management practices for Internet access services as those limitations may impact the capabilities of network operators to provide their other services. This, of course, places a great burden on the Commission, and, in this document, a group of network engineers have joined together to assist it in its task.

These engineers have developed a set of "bright-line" rules the Commission would employ should it decide to intervene in network management activities. These rules reflect the increasingly complex world of network management as users access more innovative and demanding applications and as managers need to handle rapidly surging traffic, an evolving set of applications, spam and malware, and emergency messages along with demands from a diverse array of customers with much different requirements. The following rules provide the minimum set of tools needed by network managers:

- **Should the government decide to adopt any rules concerning network management, they should apply only to the provision of Broadband Internet Access Service and not to the offering, operation, or management of any**

other services, including managed or specialized services, offered by the provider even if such services are provided over common facilities. As demonstrated earlier, modern communications networks are used to provide a wide (and constantly changing) array of services, each of which is comprised of features and functionality sought by users and each of which accordingly places different demands on the network. In determining whether and how to provide these services, network operators weigh numerous factors and make complex calculations about whether to invest substantial amounts of capital to deploy additional capacity or manage within current capacity. Further, a balance made by one entity likely is irrelevant (or at best of minimal relevance) to a balance made by another – after all, today’s communications networks vary tremendously, even among providers offering the same or similar arrays of services.

- **Broadband Internet access providers may offer users different classes of services with varying performance capabilities, which may include the imposition of limitations on users’ bandwidth and consumption subject to full and understandable disclosure.** It is a common and well-accepted practice today for broadband providers to offer classes (tiers) of service, which are sold at different prices for different performance capabilities (*e.g.* bandwidth) and with different terms and conditions (*e.g.* limitations on consumption). The key is to make sure users are given sufficiently clear information about these classes of service – both capabilities and limitations – so they can make informed decisions.

- **Broadband Internet access providers may provide quality of service or other performance guarantees to providers and users of applications so long as they are not offered on an exclusive basis for any particular application category.** Many Internet applications are highly-sensitive to latency and jitter and packet loss – and newer real-time video applications also must have sufficient bandwidth to provide a satisfactory experience for users. Providers of applications and users of these applications would benefit from Quality of Service (“QoS”) guarantees and, in fact, often request them. (Many of these providers already use Content Delivery Networks (“CDNs”) to ensure QoS from the host server to a server collocated with the broadband Internet access provider.) So long as any QoS guarantees are offered on a non-exclusive basis – that is to all providers of any particular category of application – they should be permitted.

- **Broadband Internet access providers may choose to increase or alter network capacity by building additional facilities or installing or tuning electronics -- and the government can provide incentives to build additional capacity¹ -- but providers should not be required to do so.** Decisions regarding increasing network capacity are exceedingly complex, involving not just network engineering concerns but all the factors that are part of a business case. In addition, these

¹ The FTTH Council, in fact, believes the government should be involved in providing incentives to accelerate deployment of FTTH and other next-generation access networks, and it has consistently advocated for measures to accomplish this objective, including investment tax credits and tax credit bonds.

decisions often have to be made relatively rapidly given the growth of Internet traffic. Further, the circumstances that lead to the possible investment in more capacity may be driven by a small subset of subscribers. Finally, it would be unprecedented for the government to mandate or prevent the construction of facilities and equipment by communications providers other than those who have a monopoly.

- **Broadband Internet access providers may throttle and shape Internet access traffic (1) so long as they do so for all applications of a particular type and so long as the action is of a short duration and not repeatedly imposed, (2) if it is imposed on a individual user, so long as such action is part of the terms and conditions agreed upon by the user when choosing the service, or (3) in emergency situations .** Because of the potential high volatility of Internet traffic, network managers often face surges beyond those anticipated by and engineered into the network. To ensure a satisfactory level of service, network managers should be able to address these events in several ways. First, they may throttle or shape traffic if it applies to all applications of a particular type. Second, they may throttle or shape traffic for a particular user if the provider informs the user in clear and understandable terms in advance that such activity may occur. Third, they may throttle or shape traffic if emergency situations – to permit critical traffic to get through or to prevent the network from crashing.

- **Broadband Internet access providers may permit CDNs² or others, including themselves or affiliated entities, to install and have traffic directed towards local caching equipment.** CDNs are so widespread and have such generally recognized benefits for users and content and applications providers that they should be encouraged – even though they may give advantages to particular entities (as opposed to particular applications).

- **Broadband Internet access providers may block or otherwise restrict spam, malware, and similar traffic harmful to the network and unlawful traffic (or unlawful transfer of content) – as well as provide priorities for emergency traffic and secure transmissions.** There is a virtually unanimous agreement that these activities are either so harmful in the case of spam and malware or beneficial in the case of emergency and secure traffic that network managers should be able to address them expeditiously by use of practices generally accepted in the industry.

- **Broadband Internet access providers should publicly disclose information about network usage limitations or conditions and general information about their network management practices but may limit disclosure of network**

² As will be discussed more fully in the document, CDNs have deployed vast “parallel” networks to the public Internet with the objective of shaving microseconds off transmission times and providing QoS guarantees. CDNs, like Akamai, Google, Limelight and many others, accomplish this task by transmitting content and applications from the host servers directly to the CDN’s own servers (caches) that are collocated with or located near and connected to the facilities of a broadband Internet access provider.

management practices that would negate or otherwise undermine any permissible network management practice. Providing users and applications and content providers with sufficient and understandable information about network usage limitations or conditions is important to ensure they make informed choices about subscribing to the service; however, providers should not be ordered to disclose information regarding the precise nature of any network management practices, including technical methods employed to manage traffic. Such granular detail about network management practices is often proprietary, crucial to preserve network security and reliability, and is constantly changing in response to new congestion problems or malware attacks.

II. INTRODUCTION

Hundreds of millions of users and devices have access to the nation's communications infrastructure; yet, at any given time, only a small percentage of those access the many networks making up that infrastructure. As a result, it is more economical to design networks to maximize the sharing of that infrastructure rather than dedicate significant portions of the network to specific users. Indeed, shared facilities have been a key part of the nation's communications infrastructure since its inception.

The inherent risk of a shared architecture -- or with networks with insufficiently dedicated and sized facilities -- is congestion when usage exceeds planned-for capacity during peak or near-peak usage periods, which can occur due to large number of users and devices, extremely high capacity applications, or a combination of the two. Such congestion occurred occasionally on the public switched telephone network, such as on Mother's Day, when peak capacity was insufficient to handle the surge in traffic and many users found their calls were effectively blocked. The essence of network (traffic) management, engineering the trade-off between capacity and service control so that adequate quality of service is maintained at a reasonable cost.

The traditional public switched telephone network supporting Time Division Multiplexed ("TDM") voice services and limited data offerings in a monopoly environment was far simpler than today's multi-protocol, mixed use, highly dynamic networks operated in a highly competitive environment. Nonetheless, traffic management was already a complex undertaking even for the traditional networks, given different residential and business calling patterns over various types of facilities and services at different times of day. But whatever complexity was presented by traditional networks pales in comparison to the challenges network managers and engineers face today.

Today, voice traffic may be carried over traditional TDM, virtual private, non-nomadic Internet Protocol ("IP"), or Internet-based IP ("VoIP") networks using a variety of combinations of dedicated and shared facilities -- some of which may be obtained from wholesale suppliers -- private IP networks, or the Internet. Data traffic may be sent via TDM technology or any number of packet offerings, including IP, some of which flow to and from the Internet. Video traffic may be in traditional "cable" or "Radio Frequency" ("RF") format, IPTV, or IP video streamed or sent as large files to and from the Internet. Users accessing the Internet over the same networks present a

wide variety of usage profiles and utilize myriad applications offered by third-parties unaffiliated with the network managers that place a broad range of capacity and latency demands on networks. Many local networks often operate in a wholesale capacity to other wireline and wireless providers of voice, video, and data services in addition to providing retail services. Network engineers must address rapidly changing services or applications with accompanying differences in network flows and requirements, interconnections with and hand-offs to and from different networks, constantly mutating and more malicious malware, legal requirements to protect content, law enforcement mandates, and emergency communications. Further, while congestion may occur on local networks, it is quite possible that perceived congestion and degradation in service may be determined by events and network choices of third parties that occur before content is even delivered to local networks. Moreover, end users themselves may experience congestion because of deficiencies in their own terminal equipment, inadequate computer processing power while running capacity-hungry or multiple applications, and bottlenecks within customer operated home or business networks.

Modern traffic managers, of course, can remedy, at least temporarily, some of the issues that arise on their networks by deploying additional capacity. That is the benefit of fiber-to-the-premises ("FTTP") access networks, which have tremendous capacity, reduced dependence on shared facilities, and can be readily upgraded. But, given the growth and evolution of video communications and spam and malware, even local access networks based on fiber architectures need to address the trade-off between the creation of new capacity and load management. In addition, capacity increases will address only a subset of possible issues to efficient network management. Thus, traffic management remains essential.

The Federal Communications Commission now proposes the adoption of six network neutrality rules which the agency tentatively concludes is necessary to ensure users of broadband Internet access services can access content and applications and connect devices of their choice and to prohibit providers of these services from discriminating against any content, applications or services.³ The proposed rules would

³ The proposed rules are: subject to reasonable network management, a provider of broadband Internet access service:

1. would not be allowed to prevent any of its users from sending or receiving the lawful content of the user's choice over the Internet;
2. would not be allowed to prevent any of its users from running the lawful applications or using the lawful services of the user's choice;
3. would not be allowed to prevent any of its users from connecting to and using on its network the user's choice of lawful devices that do not harm the network;
4. would not be allowed to deprive any of its users of the user's entitlement to competition among network providers, application providers, service providers, and content providers;
5. would be required to treat lawful content, applications, and services in a nondiscriminatory manner; and

allow providers to engage in “unspecified” reasonable network management practices. Thus there is a need for clear guidance as to what types of network management practices will be allowed. In addition, because broadband Internet access services are just one of the many services offered simultaneously over modern communications networks and thus the traffic patterns associated with Internet access and use are only a subset of the elements engineers must factor into overall network management calculations, it is important that the Commission tread carefully in establishing any limitations on traffic management practices for Internet access services as any limitations may impact the capabilities of network operators to provide their other services. This, of course, places a great burden on the Commission, and a group of network engineers have joined together to assist it in its task. In this primer on network management practices, these engineers seek to educate government policymakers and provide sufficiently precise benchmarks for determining what practices are reasonable.

III. BASICS OF NETWORK MANAGEMENT

A. Traditional Telephone Networks

Traditional telephone networks included both a facilities network, a series of interconnected nodes (points of aggregation) and links (*e.g.*, trunks, and loops), and traffic networks, particular interconnections of facilities for the routing of traffic which are associated with the offering of a service or internal control. The basic rules for transmission, which in turn dictate the trade-off between nodes and links, are: traffic must be carried among customers dispersed over large areas; traffic may be generated by any customers at any time to any customers; the exchange of traffic must have minimal delay. Designing an efficient network providing low-cost quality services usually means combining the proper balance of nodes and links with various traffic management techniques to handle unexpected and infrequent peak demands. Because connections between users on traditional telephone networks are dedicated end-to-end for the duration of calls, a key technique is oversubscription (or over-provisioning), which employs probability theory to enable sharing of links or ports with capacity less than that required if all users accessed the facility simultaneously.

Within traditional networks, engineering was already complex. For instance, switch-design affects transmission facility design, and decisions made in one part of the network will affect capabilities in other parts. Not only did network engineers trade-off among nodes and links, they needed to account for the capacity of facilities, how to interconnect facilities to establish circuits with sufficient redundancy, and when to employ additional electronics (*e.g.* multiplexing) to increase capacity on existing facilities, all the while assessing the cost consequences. Network engineers monitored

-
6. would be required to disclose such information concerning network management and other practices as is reasonably required for users and content, application, and service providers to enjoy the protections specified in this rulemaking.

the networks' performance closely and employed extensive modeling to examine different deployments loaded with traffic over time.

There was a great incentive for network operators to manage traditional networks effectively and efficiently once they were deployed because of the great expense of adding capacity. Simply providing additional facilities to handle infrequent dramatic peaks was not cost-effective. Consequently, operators constantly sought other solutions to deal with unexpected peak loads. Congestion could occur at any of the switching or shared transmission facilities (trunks) throughout the network. Over time, a series of network management controls were adopted which depended on detection of potentially congestive traffic patterns. These controls were either "restrictive" – because they eliminated or reduced routing alternatives so traffic did not reach the point of congestion – or "expansive" – because they used non-standard routing to bypass the point of congestion. Examples include Code Blocking (restrictive) – blocking calls according to the destination code – and Rerouting (expansive) – routing traffic to new trunk groups not normally used for the traffic in question. These techniques were largely confined to circuit-switched services such as traditional voice and now antiquated circuit switched data.

B. Management of Modern Communications Networks

In the past twenty-five years, packet-based and eventually IP-based services have burgeoned, and networks have become more capable and more complex. Telephone providers have deployed FTTP and higher-capacity DSL networks and are frequently managing "cable" video or IPTV traffic. Cable operators too found themselves increasing the nature and scope of their traffic to include voice and broadband Internet access. Most networks used to provide service to retail customers also serve wholesale customers.

In packet-switched networks, data are divided into packets with a specific format (including destination and control information) and length. Unlike circuit-switched networks, there is no end-to-end path dedicated to a call for its duration in such networks. Rather, in packet transmissions, the packets are disaggregated and generally sent over a variety of routes to the destination where, using control information, they are reassembled to make a coherent message. When the load on a route gets too great, packets are placed in a queue (memory buffers) until they can get transmitted. This means that packets can be received at the destination out of order, but the control information ensures they are reassembled properly. However, that does not mean the quality of the transmission is satisfactory for all services. Too much variation in the required time to transmit a packet (the propagation delay or "latency") leads to greater "jitter," which is unacceptable for VoIP calls. Excessive jitter also can be caused by substantially increased loads leading to packets being discarded (packet loss). Packet transmissions thus have different sensitivities to transmission parameters, such as transaction rate and length, and different management requirements.

A typical modern local access network can provide a variety of services – both circuit and packet-based, including:

- Voice services which may either be circuit (TDM) based or packet-based VoIP;
- Broadband Internet Access Services which may include caching for content delivery networks;
- Video services which may be RF or IPTV;
- Carrier-grade services for business customers to multiple locations(e.g. virtual private line);
- Wireless backhaul services for mobile telephony providers;
- Wholesale “lit” capacity to other carriers.

Each of these services is distinct, although they may be sold in bundles, and each has different network requirements in terms of bandwidth, latency and jitter, and other parameters to ensure satisfactory quality for the user. For instance:

- Voice services often require very low latency, jitter, and packet-loss, in addition to end-to-end timing synchronization.
- Many services dependent upon Internet access, such as e-mail and web searches, can tolerate more jitter and data latency when real time transmissions are less critical.
- Carrier-grade business services are usually provided with Service Level Agreements dictating latency, jitter, packet-loss, and other parameters.
- Streaming video requires significant bandwidth, especially for high-definition signals, and low latency and jitter, and very low packet-loss.
- Mobile backhaul services require extremely low levels of latency and jitter and timing synchronization.

Only one of the services listed above, Broadband Internet Access Service, connects to the public Internet and is the subject of the FCC’s proceeding. Yet, for most providers, all of these services share many of the same facilities, and they need to be managed not only individually but collectively. Today, for example, TDM voice traffic is usually given priority over broadband Internet access traffic to avoid problems caused by bursts in Internet transmission. If voice services were susceptible to congestion caused by such bursts, calls could be rendered unintelligible or be dropped. This practice does not “harm” Internet access traffic because voice traffic levels are, generally, predictable, and IP packets can be queued or regenerated when voice demand rises.

Cross-service network management also occurs between IPTV services and Internet access services. When these two services traverse the same access network, the network typically allocates bandwidth for both services to be provided

simultaneously. Otherwise, large bursts of Internet access traffic can degrade IPTV picture quality or even block the video signals altogether.

Practices of bandwidth or channel allocation – either physical or virtual – among different services often occurs in response to customer requirements, particularly for business customers who require more than “best efforts” service or for public safety entities who need priority for emergency communications. Few businesses can afford to have their own networks, yet they require high-quality, “uninterruptible” service. The most efficient way to achieve this objective is by carving out capacity in shared network facilities, that is, creating “managed” services and guaranteeing “Quality of Service” (“QoS”) through Service Level Agreements (“SLAs”).

Network management also is required for any single service. Traffic over broadband Internet access services can place enormous, dynamic stresses on networks at the same time consumers and applications providers are demanding a higher degree of QoS in constantly evolving and shifting patterns. This presents network engineers and managers with substantial, often shifting, challenges. As discussed below, they need to have the tools to respond promptly and the certainty that their actions are legal.

In sum, for modern communications network providers, network engineering and management is a serious business. Customers complain to their providers promptly and continuously if service quality deteriorates – even if, as discussed below, it is not the fault of the provider. To further ensure they act responsibly, network providers, in conjunction with their industry groups, have spent enormous amounts of time developing elaborate network architecture and framework requirements in response to growing and shifting user demand. The telecommunications industry has established requirements for FTTH (ITU) and DSL (DSL Forum) networks, and for the cable industry, DOCSIS (Cable Labs) network requirements. These standards, which are discussed and propounded in an open process, provide key benchmarks for anyone judging the reasonableness of network engineering and management actions.

IV. NETWORK MANAGEMENT OF BROADBAND INTERNET ACCESS SERVICES

A. Overview of Internet Traffic

1. Introduction

a. Transmission Control Protocol

Transmission Control Protocol (“TCP”), the most widely used protocol in the transport layer on the Internet (*e.g.*, HTTP, TELNET, and SMTP), plays an integral role in determining overall network performance. Fundamental to TCP’s design is that the end stations involved in a communication (*e.g.*, a web site hosting server and a computer connected to the Internet) are responsible for controlling data flow between them. There are not explicit signaling mechanisms in the network which tell the end stations how fast to transmit, when to transmit, when to speed up or to slow down. The

TCP software in each of the end stations controls the flow from implicit knowledge it obtains from the network or the explicit knowledge it receives from the other end station.

One of TCP's primary functions is to properly match the transmission rate of the sender to that of the receiver without overwhelming the network. TCP must rely mostly upon implicit signals it learns from the network and the remote end station. TCP must make an educated guess as to the state of the network and trust the information from the remote end station in order to control the rate of data flow.

TCP includes congestion control techniques, which are defined in the Internet Engineering Task Force ("IETF") standard, RFC 2581. This reference document specifies four standard congestion control algorithms that are in common use: Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery.

Slow Start is a mechanism used by the sender to control the transmission rate. This is accomplished through the return rate of acknowledgements from the receiver. During the initial data transfer phase of a TCP connection the Slow Start algorithm is used. However, there may be a point during Slow Start that the network is forced to drop one or more of these initial packets due to overload or congestion. If this happens, Congestion Avoidance is used to slow the transmission rate. Slow Start is used in conjunction with Congestion Avoidance as the means to manage data transfer at what appears to be the highest available speed at the time. This mechanism will force the sender to more slowly grow its transmission rate as it approaches the point where congestion had previously been detected. The Fast Retransmit algorithm is used when duplicate acknowledgments are received. This provides implicit knowledge to the TCP sender that data is still flowing to the receiver and gives a strong indication that serious network congestion may not exist. Instead of reducing the flow of data abruptly by going into Slow Start, the Fast Recovery algorithm allows the sender to resume transmission at a higher throughput. Although RFC 2581 and its associated algorithms have performed well in ensuring top performance on TCP/IP networks, much work has gone into enhancing TCP performance and responsiveness to congestion even further.

b. The Internet as a "Best Efforts" Network

The Internet transmits packets in no particular order, on no fixed path, and on a "best efforts" principle with no assurance they will reach their destination. Dropped packets may occur, for instance, due to network congestion, when routers run out of buffer memory. When edge computers reassemble these packets coming from different routers into the original message, they are able to recognize if packets are missing and, if so, ask for them to be resent. Consequently, the Internet is inherently "friendlier" to applications that can tolerate the delays involved when dropped packets are recognized and then resent.

Viewing a static web page or one with a small amount of animation imposes minimal traffic requirements. Actively downloading or uploading a file causes a burst in traffic. If there are a sufficient number of users engaging in similar activities at once, it can lead to congestion and delay. The number of users that triggers this issue

depends upon the type of application(s) and overall network capacity. At its most harmful, congestion can temporarily halt the transmission of traffic on key links or large parts of the network. This could occur because of a denial of service attack. It also might be driven by certain types of technologies or applications. For instance, as discussed later, versions of P2P technologies, which often transmit large file sizes, can consume substantial amounts of bandwidth on many links and pose a real problem for platform providers. There also are applications that are not “TCP-friendly” because when congestion occurs they do not reduce their rate of transmission. Even some lesser instances of congestion can lead to jitter. Jitter has little effect on most browsing activities, but it can noticeably affect applications that require steady, interactive, real-time transmission, such as on-line gaming and VoIP.

There are serious issues in addition to congestion that arise in the transmission of Internet traffic and which must be managed. Malware attacks are frequent. While most of these are annoying, some can be devastating – and virtually all require attention from network managers. Just-in-time applications, such as emergency services or VoIP transmissions, often suffer because of traffic overloads or equipment failures. There are also grave concerns because the network does not guarantee secure transmissions.

2. Types of Traffic and Their Requirements

a. Beneficial Traffic

Internet delivered services and their network requirements vary tremendously. Traditional services like email (often called bulk applications) generally require minimal bandwidth and can tolerate significant latency and jitter and packet loss/retransmission. Transmission of large-size files (*e.g.* video files) for later use require significant bandwidth but too are less sensitive to latency and jitter and packet loss. Web-surfing also historically has been a bulk application. But as more sites include more interactive features, this activity requires increasingly greater bandwidth and lower latency and brings into question how much longer it will squarely remain a bulk application, if it still is one today. In general, for bulk applications, the Internet’s “best efforts” practice may be sufficient. Further, enhancements, such as local caching of websites, can improve performance. In contrast, as noted above, VoIP services, while requiring only limited bandwidth (voice packets use little capacity), need low latency and jitter – because these calls tend to be highly interactive. Video streaming (and other paced and burst-paced applications) not only needs low latency and jitter but high bandwidth as well – and HD and 3D video will require significantly more.

For most traditional applications, network engineers only need to factor in downstream bandwidth requirements. However, that is not the case for more recent and future video-based applications. Peer-to-Peer (“P2P”) applications, which can support either large-file or streaming video transmissions, require greater upstream bandwidth. This also is the case for services such as and two-way video communications (*e.g.* telepresence), which are becoming more common.

In addition to the applications and considerations just described, certain applications are of particular concern for network managers:

P2P Applications – P2P networks are built by individual users downloading the application, enabling each to become a participant in the network to share information. In pure (unstructured) P2P networks, users (nodes) act both as clients and as servers, that is, all are equals (peers) in receiving and sending data. When a user wants data, it sends a signal to as many peers as possible – each of which may respond, assuming it has the requested data, by sending all or just a portion of the data. For purposes of traffic management, because of the constant signaling and data transmissions involved with P2P networks, traffic flows tend to be very large but are not optimized, placing disproportionate loading on the network. In 2008, P2P comprised about 50% of consumer Internet traffic. P2P providers have sought to alleviate some of this load by altering the pure network structure to include some greater central control. They also are working with broadband providers through the P4P initiative to find ways to better optimize flows.

Gaming Applications – Gamers are frequent and demanding users of Internet access services, and the number of users participating in multi-party, on-line gaming is expected at least double in the next several years. Because gamers demand rapid response capabilities, services must have low latency and sufficient bandwidth. In addition, because gamers play over an extended period, there is a “longer peak” during which these networks must meet these performance requirements. Game designers undoubtedly are able to play a role parallel to those of network managers by increasing playability while minimizing demands placed on networks.

Cloud Computing and Associated Applications – With cloud computing, processing that once took place at users’ computers now occurs in the Internet “cloud.” For instance, instead of having a word processing or spreadsheet program on a user’s computer, the program is on a server in the “cloud” and a user accesses it via a web-browser. Such an application can save users money, permit easy updates, provide back-up of data, and limit infections. At the same time, it generates considerably more traffic and places greater demands on Internet access networks, both downstream and upstream. In addition, because cloud computing can work with any number of applications, with more demanding applications like fast-action video gaming, the demands on the network increase dramatically. Such demanding activities require greater amounts of bandwidth and low latency and jitter and packet loss.

b. Malicious traffic

In addition to the many beneficial applications, the Internet has spawned a great array of harmful traffic, which network managers must confront throughout every day. Malware, viruses, botnets and other malicious programs pose continuous threats to networks and user equipment. The website, *Viruslist*, provides regular reports on this threat. The following statistics on malware are from November, 2009:

Malicious programs on the Internet

Name	Number of attempted downloads
Trojan-Downloader.JS.Gumblar.x	1714509
Trojan-Downloader.HTML.IFrame.sz	189881
Trojan-Clicker.JS.Iframe.be	170319
<u>not-a-virus:AdWare.Win32.Boran.z</u>	136748
Trojan.JS.Redirector.l	130271

November trends

The overall picture remained unchanged in November. At the moment, the most common strategy for spreading malware is to use a malicious script + exploit + executable file. More often than not, this is how malware designed to steal confidential data or extort money from users is spread... Another marked trend of recent months that continued in November was the use of websites created using standardized templates to spread rogue antivirus solutions. Cybercriminals are also aggressively using packers (usually polymorphic) in the hope that this will help the packed malicious programs avoid detection, so they won't have to make significant modifications to the malicious programs themselves. This month malware was also distributed via P2P networks using multimedia downloader programs, a method that the cybercriminals made use of last December.

For purposes of traffic management, it is important to note that when malware, such as a botnet program, infects a user's computer, it can result in greatly increased traffic flow as the malware seeks to infect or otherwise harm networks, servers, and other computers.

Spam -- which comprises most email traffic -- also continues to be a major nuisance for traffic managers. Again, *Viruslist* has produced recent statistics:

Spam evolution: October 2009

Recent trends

The amount of spam in email traffic decreased by 0.6% when compared to September's figure. The total average being 85.7% for October.

Links to phishing sites were found in 0.9% of all emails, an increase of 0.1% when compared to September.

Malicious files were found in almost 2% of all emails, an increase of 0.7% when compared with the previous month.

Halloween and Christmas themes were actively exploited by spammers.

Spam in mail traffic

The amount of spam detected in email traffic averaged 85.7% [of total email traffic] in October 2009. A low of 81.2% was recorded on 3 October with a peak value of 89.7% being reached on 18 October.

c. Legal Requirements

In addition to the issues just discussed about types of Internet traffic, broadband Internet access providers need to account for other factors when managing their networks. Broadband providers have various legal obligations, including ensuring compliance with Communications Assistance for Law Enforcement Act ("CALEA") – which require them to monitor traffic for law enforcement purposes – and the Digital Millennium Copyright Act ("DMCA") – which requires them to remove pirated material from transmission. They also must work with parents to control access by children to objectionable content.

3. Traffic Growth and Evolution

The growth and evolution of Internet-based applications and their demand for greater amounts of bandwidth and more rigorous performance requirements is almost overwhelming. Napster (and the great surge of P2P traffic) began its meteoric rise just ten years ago. While its fall only came a few years later, P2P has lived on through BitTorrent and many other providers. Apple's iTunes store and Skype's VoIP service launched in 2003, YouTube in 2005, and NBC's Hulu in 2007. In the next few years, we expect users to regularly upload HD video and swap 3D streams. Cisco's most recent report, *Hyperconnectivity and the Approach Zettabyte Era* (June, 2009) sums up this inexorable growth in traffic, especially the increase in video transmissions:

- Internet video is now approximately one-third of all consumer Internet traffic, not including the amount of video exchanged through P2P file sharing.
- The sum of all forms of video (TV, video on demand, Internet, and P2P) will account for over 91 percent of all global consumer traffic by 2013.
- Video communications traffic growth is accelerating.
- Real-time video is growing in importance.
- Video-on-demand (VoD) traffic will double every two years through 2013.

As network engineers for local broadband providers, we can confirm this trend. For Internet access services, traffic is doubling approximately every six to twelve months. Judging from the most recent activities by users to stream video directly from websites, we see no let up in the near term. As a result, every entity or person in the Internet eco-system – from those operating content servers on one end to end users on the other end -- is upgrading capabilities to support these new applications and demands. Increasing capacity and capabilities, however, does not occur instantaneously in each segment of the eco-system nor may it be deemed cost-effective. In addition, in the face of these efforts, applications and content providers keep churning out newer, more demanding versions of their products. Consequently, congestion – or the threat of congestion -- is a persistent fact of Internet life and will remain so.

B. Points and Sources of Congestion for Internet Traffic

There are various parts of the Internet where aggregation and interconnection occur or where information is processed and transmitted. Each of these can be the source or locus of congestion (or a diminution in service quality).

User Equipment and Networks – Computers, local area networks, and modems are all potential sources of congestion. Because malware is an ever-present threat and computers are often insufficiently protected, many home computers are infected – possibly as many 20% at any time. If they have been captured by botnets or are part of “zombie” networks, infected computers may constantly transmit information – potentially millions of emails every day -- as they search out targets. An infection thus not only slows a computer’s performance. In the aggregate, infected computers can generate enormous amounts of Internet traffic.

Traffic within a premises also may be degraded or blocked because of inadequate or defective modems or local area networks (e.g. an out-of-date Wi-Fi network with insufficient bandwidth or interference received from household appliances). While these problems may be isolated to the premises and not affect a broadband provider’s network or the larger Internet, the user experiences a diminution in quality of service sometimes indistinguishable (to the user) from problems occurring caused elsewhere in the network. Network providers frequently handle complaints from users experiencing such “home-grown” difficulties.

Local Access Networks (including “2nd mile” shared facilities) – Wireline local access networks include connections from the premises to the first point of switching/routing (or to the “central office” or “headend”). The potential for congestion varies largely based on the capability of the technology and the degree of sharing. In certain instances, such as with active FTTP networks, passive FTTP networks where the “splitter” is at a central office, and home-run copper DSL networks, facilities are completely dedicated to each premises. Where facilities are dedicated (and not shared), it is easier to ensure each premises always achieves its

defined performance level for the service. (For DSL (copper-based) networks, the level of performance decreases significantly as distance from the router to the premises increases. With FTTP, which has much greater capacity than DSL networks, distance is not critical.) However, in most instances, facilities are shared among users beyond the first point of switching/routing, and the degree to which it is shared and the overall capability of the network will determine the performance level a user receives, especially at times of peak usage. For instance, with passive FTTP networks where the splitter is placed in the field, capacity may be shared among 32 or fewer users, but because overall network capability is so great, the effect on performance is minimal. In contrast, with cable hybrid fiber-coax ("HFC") networks, sharing is more extensive and bandwidth is more limited. Hence, especially with HFC networks that have not been upgraded to DOCSIS 3.0 technology, the opportunities for congestion and degraded performance are greater, particularly at periods of peak usage. Fiber-to-the-Node ("FTTN") networks, where fiber is run to a neighborhood but copper is still used to connect to the premises, face similar issues.

Middle Mile Facilities – Middle mile transmission facilities connect the local access network to the Internet node. For broadband Internet access services, these facilities are shared among all users of a given local provider and, consequently, have substantial capacity. The problem is that with two-way video-related traffic growing so quickly, whenever new capacity is added, the middle mile may be filled quickly. As noted above, broadband providers in general are finding that traffic is doubling at least every two years – and for some individual providers the rate of doubling is even faster, every six to twelve months -- which means frequent additional investment in middle mile capacity.

Internet Backbone Facilities – While Internet backbone traffic was initially provided by government agencies, for many years, commercial providers built and operated the infrastructure from Internet access nodes to servers originating content and applications. Over the years, as Internet traffic has grown, these backbone facilities have deployed enormous "pipes." The problem is that as traffic growth accelerates, current capacity needs to continue to increase. As a result of the potential mismatch between traffic and capacity on these backbone facilities, congestion may occur from time-to-time.

Content and Applications Servers -- Servers operated by content and application providers are sophisticated computers storing and transmitting data. As such, they require sufficient processing capacity. Like other broadband Internet facilities, many servers have become overwhelmed by the growth of video traffic and have acted as a bottleneck which is experienced by the end users in a way indistinguishable from when they experience congestion or delays occurring in their local providers'

networks. In other words, even if a user has sufficient bandwidth from its broadband service provider to receive a streamed video, even if there is sufficient shared local network facilities, even if there is adequate capacity in the middle mile, and even if uncongested Internet backbone facilities are available, a user may not receive the video with sufficient quality. Many application providers have sought to address congestion problems at their servers, as well leapfrog others that may occur in backbone networks, by using content delivery networks (“CDNs”). CDNs provide direct connections between the main servers for websites and additional (mirror or local caching) servers that are collocated with broadband access providers. When CDNs are employed, much of the traffic between a user and a website does not need to traverse the entire Internet but instead can be handled just by the local caching server.

C. Network Management Techniques Employed Today by Broadband Internet Access Providers

1. Introduction: Users and Providers Address Congestion

As seen in the previous section, congestion in the Internet can occur at many places not under the control of the broadband provider. Most users, however, have no understanding of the underlying cause of the congestion they experience and turn to the Internet access provider whenever broadband service quality is degraded. As a result, providers have set up extensive, round-the-clock customer and technical service operations. Some even have established “geek squads” that go to customer’s homes to determine the source of the problem. To facilitate these operations, providers have installed elaborate monitoring and detection equipment to understand more precisely traffic flows and points of congestion. Providers have learned the value of giving subscribers understandable and complete information – in other words, educating them – since it tends to reduce future issues and lowers churn. In essence, in almost all instances, there is a close and positive working relationship between subscribers and providers in addressing concerns raised by congestion.

The following sections describe tools used by broadband Internet access providers to manage their networks. These tools may be used alone but more often are used in combination. In addition, they frequently evolve to respond to new concerns and new technologies.

2. Increase Network Capacity

In seeking to ensure the provision of high-quality service, network providers continuously upgrade the capabilities of their infrastructure – which, for major upgrades, is a costly and non-trivial undertaking. The need to deploy additional infrastructure has become even more acute for broadband Internet access providers because of the explosion in video traffic. Each year, local network providers spend many billions of dollars to deploy new infrastructure. In fact, most recently, there has been a

concerted effort by access providers – as well as others in the Internet chain -- to expend capital to dramatically increase capacity. Most of these efforts involve driving fiber closer to the premises because of its greater performance characteristics. But, capacity also can be increased by upgrading electronics or bonding channels.

Verizon, for instance, will soon complete its deployment of FTTP (FiOS) to 18 million households, and AT&T will deploy FTTN, with nodes approximately 3000 feet from premises, to 30 million households. The cable industry should complete its deployment of DOCSIS 3.0 plant to about 100 million households within the next several years and will continue to push fiber deeper into the network to serve nodes with fewer users. The following chart provides an overview of the greatly increased broadband speeds provided by each of these technologies today and those expected over the next five years:

Local Broadband Access Network Technologies

Technology	HFC		FTTH		FTTN / VDSL2	
Transport Mechanism	DOCSIS 3.0 8-Bonded Down 4-Bonded Up		10G-PON 25G/125G		3000' Vectoring 1000' Non-Vectoring	
Transmit direction	Down	Up	Down	Up	Down	Up
Line Rate (Mb/s)	152	9	250	125	35	10
Homes/Node (split ratio)	250		32		1	
Penetration	35%		100% ²		NA	
Subs/Node	88		32		NA	
Peak % Subs Online	40%		40%		NA	
% Active Tx/Rx	20%		20%		NA	
Available Mb/s Sub	22	15	975	480	35	10
AVAILABLE TODAY						

Technology	HFC		FTTH		FTTN / VDSL2	
Transport Mechanism	DOCSIS 3.0 8-Bonded Down 4-Bonded Up		10G-PON 25G/125G		3000' Vectoring 1000' Non-Vectoring	
Transmit direction	Down	Up	Down	Up	Down	Up
Line Rate (Mb/s)	304	108	10000	250	100	28
Homes/Node (split ratio)	125		64		1	
Penetration	40%		100%		NA	
Subs/Node	50		64		NA	
Peak % Subs Online	40%		40%		NA	
% Active Tx/Rx	20%		20%		NA	
Available Mb/s Sub	76	27	1950	480	100	28
DEPLOYED BY 2015						

Current Generation Access

Next Generation Access

- Cablevision currently offers 101Mb/s down 15Mb/s up as a top tier
- Assumes splitters (and associated OLT cards) fit in series; result is worst case scenario for FTTH
- Using same statistical approach as HFC; FTTH carriers can offer 975Mb/s but current offers do not exceed 100Mb/s
- Voice, Video and Data share bandwidth on the same VDSL line
- DOCSIS 3.0, 8-channel bonded downstream and 4-channel bonded upstream equipment will be available in 2010
- 10G-PON standard --802.3.187.2 completed in September 2009
- Vectoring expected to be approved in the ITU G.993.5 standard May 2010
- Cioffi, FCC Presentation November, 2009

Other sources: Motorola, Cisco, SNL Kagan, Corning, CSMG

In examining this chart, the focus should not be just on the ultimate speeds but on factors upon which such speeds are dependent: technology used, degree of sharing, and expected peak usage. All of these enter into the calculations of network engineers and determine what additional network management will be required. These factors also enter into the overall business model and determine whether deployments are cost-efficient.

Broadband Internet access providers also have been adding capacity to their middle mile links. Again, because of the growth in traffic, for most providers middle mile capacity has been doubling every two years (if not more frequently) recently as video traffic has exploded. As a consequence, the cost per megabyte of information delivered over these middle mile links has dropped precipitously – by approximately 60%.

3. Establish Classes (Tiers) of Service for Users

Broadband Internet access service is typically provided in tiers, each with a different set of performance capabilities (which may be actually delivered or advertised or “up to” throughput capabilities for downstream and upstream throughput) from the premises to the Internet access node. It also is possible to add other performance criteria – maximum latency and jitter and packet loss – to distinguish tiers. Each tier of service is set at a different price, which increases as the speeds or other capabilities increase.

Network engineers take steps to maximize the potential that the shared facilities are able to meet the performance levels for the different tiers (in addition to ensuring sufficient dedicated capacity is provided). The various broadband tiers are generally provided on a “best efforts” basis but some can be accompanied by QoS assurances. It also is possible to offer lower tier service on a “preemptible” basis where, during periods of congestion, service is throttled (degraded) or blocked entirely.

As an additional network engineering and management technique, broadband providers may include in a tier limits on the total amount of usage over a period. This technique, which is enforced through metering, is particularly directed at the small percentage of subscribers who are high-volume users, frequently receiving or sending great amounts of data. “Well-above-the-average” activity places great stress on shared facilities and can cause congestion unless additional facilities are constructed, in which event these high-volume users would effectively make most other users pick up network costs they do not cause.

Finally, another key feature of tiering is transparency and education. Broadband providers have an obligation to provide users with complete and understandable information about each tier’s prices, terms, and conditions, including limitations on service or priority relative to other tiers. In addition, providers need to be responsive to subscriber inquiries once service is initiated.

4. Direct Management of Traffic

a. Oversubscription of Shared Facilities

Oversubscription (over-provisioning) means assigning a total committed amount of traffic – that is the traffic that would occur if all users were transmitting simultaneously -- to a given transmission link or port that is greater than that link or port’s capacity. Network engineers base the degree of oversubscription on the concept that it is extremely unlikely that all users will seek to access the network at the same time. The

upside to this approach is that networks are used more efficiently. The downside is congestion if the calculations are wrong.

Network engineers increase or decrease the probability that congestion will occur by altering the degree of oversubscription. The “right” oversubscription ratio for a network varies widely and depends on a variety of factors, including underlying network topology and architecture, the volume and variability of traffic or traffic flows, the number of users using it and subscription trends, the kind of users using it, the mix of applications running on it, historic and anticipated variation in traffic loads, and the link owner’s technology and marketing values, among other things. In other words, determining the degree of oversubscription is not so straightforward, but, where traffic patterns are more predictable – such as with traditional voice traffic – oversubscription is a useful mechanism to manage traffic. However, in the Internet environment – with frequent and continuing emergence of applications that have particularly demanding characteristics – oversubscription as a traffic management tool is much more challenging.

Real-time applications (e.g. VoIP) do not work well in the presence of latency and jitter and packet loss. P2P applications, meanwhile, can result in very large and variable increases in demand for bandwidth. Because of these and other “demanding” applications, network engineers either need to lower the oversubscription ratio, resulting in unsustainably high bandwidth costs, or face frequent congestion and unacceptable levels of service. Consequently, while oversubscription continues to be employed, other management approaches may be more cost effective.

b. Traffic Control Techniques

(1) General Traffic Control Techniques

Broadband Internet access providers have developed a series of techniques that they apply directly to address potential congestion [caused by extensive malicious traffic]:

- Blocking – This is the most restrictive network management practice and is generally reserved to stop spam or malware but also could be employed to ensure emergency traffic flows properly.
- Throttling – This practice controls the amount of traffic flowing into a network in a specific period, buffering (storing) the packets or if necessary dropping packets.
- Shaping – complex set of techniques which can control such things as the volume of traffic and the rate at which it is flowing.

The ability to apply these controls on an application-specific basis has been possible for many years using technologies like diffServ (discussed below) but has become more broadly useful as Internet traffic grows in type and volume. In particular,

techniques such as deep packet inspection (“DPI”) extend the ability to recognize applications and protocols, while policy management creates a structured framework for applying a wide range of policies to subscribers, applications and network flows.

In simple terms, in addition to establishing tiers of service as discussed above, service providers can control traffic in the network as whole or by service or application:

- Network – Network managers may control congestion or potential by throttling (limiting bandwidth) when traffic grows beyond designated limits, with no reference to the underlying application or the source/destination of the subscriber.
- Service or Application – If network managers believe that congestion is due or potential will be due to a specific application, they can either limit use of that application or give other applications higher priority.

Thus, traffic management can be dependent on or independent of the application in use, or the individual subscriber using it. Moreover, the relatively sophisticated equipment now available allows these previously separate capabilities to be combined, so that individual subscriber usage of a specific application is controlled, based on the type of subscription or tier or other information.

(2) DiffServ and MPLS

Several developments have emerged beyond TCP and oversubscription that are widely used to help manage traffic at the level of applications. The evolution of the Internet from a purely academic network to a commercial network used by enterprises and others to deliver services with critical technical needs led, for example, to the development of two IETF standards called diffServ and Multiprotocol Label Switching (MPLS).

DiffServ is a relatively simple scheme for classifying traffic into a small number of service or applications classes in order to give priority to certain traffic types. A typical diffServ classification distinguishes between VoIP, time-sensitive transaction-oriented traffic, and best-effort traffic (e.g., email). DiffServ is often used in association with MPLS, which adds a label to packets traversing an IP network. With MPLS, data packets are assigned labels and packet-forwarding decisions are solely on the contents of this label, without the need to examine the packet itself allowing the creation of virtual end-to-end circuits across the network. Among other things, these methods permit different Classes of Service (“CoS”) to be applied to different applications, so that certain applications (such as VoIP) get priority if a link is congested.

Both diffServ and MPLS were created about ten years ago and are used by enterprises and broadband Internet access providers to provide greater quality assurances for certain services. Variants include MPLS-TE (traffic engineering) and T-MPLS, which is specifically designed to route IP services like IP-telephony.

(3) Malware Detection and Blocking

To address the growing problems with spam and malware, broadband Internet access providers are installing equipment in their network that detect and block “bad” computers or websites that are sources of these harmful products. One such product is Adaptive DarkNet™ offered by MainNerve. This product resides in the provider’s network and constantly monitors transmissions from users’ computers to determine whether they are infected with malware. (Main Nerve guarantees no false positives.) It then relays this information to the broadband provider which can block transmissions from the user’s computer to the known malware IP address (prevents the infected computer from “phoning home”).

(4) Deep Packet Inspection

Internet packets contain not only the communications content but also other information that identifies where the packet came from and where it is going to, among other things. DPI as a concept first emerged around 2000, but it is not yet a standardized technique. At one level, inspection of packets can be said to be deep if it achieves the basic objective of recognizing the underlying application that the packet is carrying. How it does this is usually proprietary and often confidential and is a source of differentiation among vendors. However, the term DPI is often related to the seven layer Open Systems Interconnection (“OSI”) model for communications, which divides the task of interconnecting computer systems into seven layers.

DPI equipment inspects the contents of packets traveling across an IP network. It can identify the application or protocol in use by examining the source and destination IP address, the port number, and the packet payload. Port numbers are a basic means for identifying applications. For example, email using the Simple Message Transfer Protocol (“SMTP”) uses port 25. Packet headers include this information, along with source and destination addresses and other data, including diffServ class information where relevant. Equipment may also look for telltale signs of an application, such as the length of the packet payload. Compiling this information, DPI equipment can identify applications with varying levels of accuracy.

DPI was originally intended – and is still used – as an offline traffic monitoring and planning tool which analyzes traffic to help service providers understand what applications are consuming bandwidth and how that is changing. Today, DPI also is used to identify and address potential traffic overloads by controlling the bandwidth available to certain applications. Finally, providers use DPI to identify and block malicious applications and security threats; improve the performance of critical applications; apply parental controls on a subscriber by subscriber basis; and provide of tiered services.

(5) Policy Control

Policy control and management is a broader set of techniques than DPI that applies controls to Internet traffic flows within a structured and standardized architecture. It is not yet as widely deployed as DPI. Policy tools are in some ways competitive with DPI equipment and in others complementary. Policy control is necessarily a broad concept because it is usually based on the use of an automated rules engine to apply simple logical rules which, when concatenated, can enable relatively complex policies to be triggered in response to information received from networks, customers and applications. The set of conditions can be extended by simply adding other terms. Dynamic information (*e.g.*, the customer's location, the device in use, or network conditions at the time) can be added to the rules invoked in a particular case.

An important feature of most policy tools is that they have links into subscriber and billing databases, making policy equipment a potentially valuable means for providing more customized services to customers. It also means that policies are often related to subscribers rather than applications.

Policy architectures have been standardized over the past 3-4 years by several international telecommunications standards organizations. These include 3GPP and EGPP2, which create standards for cellular mobile network operators; ETSI, which created a policy architecture for wireline telephone companies as part of its TISPAN architecture for next-generation telecommunications networks; and CableLabs, which created a policy architecture for cable multiservice operators (cable MSOs).

These architectures typically envisage two elements – a Policy Decision Point (PDP), which is usually a highly intelligent and compute-intensive centralized device that is usually associated with a policy rules regime, and the Policy Enforcement Point (PEP). The job of the PDP is to make policy decisions on behalf of less intelligent devices. Because rules engines are generic and flexible entities, policies can be invoked to handle an indefinitely wide range of conditions, triggered by the presence of a particular application, a particular subscriber, a destination URL, or indeed any data point that can be identified and effectively acted upon.

In a basic generic architecture, policies are enforced by PEPs, which unlike PDPs are usually distributed devices closer to the subscriber. A PEP is essentially any piece of subscriber equipment that is capable of enforcing a policy decision – a wide range of equipment that includes broadband-remote access servers, gateway GPRS support nodes, DPI appliances, media gateways, session border controls, and so on. These PEPs vary in the types of enforcement they can perform, from relatively basic to relatively complex, but the basic principle is always the same.

Policy standards also include interfaces to equipment such as the Home Location Register (HLR) or Home Subscriber Server (HSS) used in 3G mobile networks. These are repositories of subscriber data, and by referring to subscriber data, policy servers can make policy decisions which directly relate to individual users.

Policy tools are more subscriber-centric than DPI tools. One noteworthy development in this area is Resource Admission Control (RAC), a concept defined in the ETSI TISPAN. Its main purpose in the standard is to allow telephone companies that are implementing all-IP networks to offer a telephone service that emulates orthodox telephone services when calls are set up – that is, when a request to make a call is received, it can deny access with a busy signal if the network is considered to be too congested to carry the call.

c. Content Delivery Networks and Local Caching

Unlike the previously discussed traffic management techniques that are solely or largely controlled by broadband Internet access providers, CDNs are developed by other providers who then reach agreement with broadband providers to locate and connect their facilities. CDNs have deployed vast “parallel” networks to the public Internet with the objective of shaving microseconds off transmission times, bypassing points of congestion, and providing QoS guarantees. CDNs, like Akamai, Google, Limelight, and many others, accomplish this task by transmitting content and applications from the host servers directly to the CDN’s own servers (caches) that are collocated with or located near and connected to the facilities of a broadband Internet access provider. These “local” servers contain sophisticated – and constantly updated – algorithms that dictate how to route information from host servers to end-users circumventing much of the traditional infrastructure of the Internet, especially backbone but, possibly, also Middle Mile. For instance, if a large group of users flood a web site seeking particular information, CDNs recognize the demand and mirror the information, placing a copy on servers closest to the users. A similar practice occurs when users demand very large files, such as popular movie downloads. CDNs also recognize breaks in transmission paths and set up alternative routes. Because they reduce opportunities for congestion (both on a broadband provider’s Middle Mile facilities but also elsewhere in the Internet), the appeal of such networks to network managers is readily apparent.

CDNs have become highly attractive to small and large content and applications entities. For a smaller entity, a CDN enables the entity to take advantage of the CDN’s scale and network reach, ensuring quality of service for the delivery of content. The arrangement also enhances network security and provides redundancy. In addition, it ensures the smaller entity could provide more than a “best efforts” service. Finally, the CDN’s significant market presence provides it with the capability to handle network problems experienced on the Internet, including any issues with broadband Internet access service providers.

It is difficult to underestimate the important role that CDNs, like Akamai and Google, play in the transmission of information via the Internet. For example, approximately 20% of all web traffic is sent each day over Akamai’s network and its 56,000 servers around the world. Akamai’s CEO stated that by improving the transmission of information over the Internet, the company is creating “emancipated consumers” that are able to access any information any time and incumbent network

providers will not be able to hinder access: “They’ll go right around the cable companies if they want.”

V. REASONABLE NETWORK MANAGEMENT PRACTICES FOR WIRELINE NETWORKS

From the forgoing discussion, it is clear that the task of network managers has become increasingly complex as users access more innovative and demanding applications. A “day in the life” of any of these network managers is filled with constant challenges as they need to handle traffic surges, an evolving set of applications, spam and malware, and emergency messages along with demands from a diverse array of customers with much different requirements. In many instances, they need to make snap decisions to prevent networks from crashing. Therefore, should the government wish to intervene in their activities, they require “bright-line” rules of the road to enable them to perform critical tasks in the time required. The following rules provide the minimum set of tools needed by network managers overseeing wireline networks:

- **Should the government decide to adopt any rules concerning network management, they should apply only to the provision of Broadband Internet Access Service and not to the offering, operation, or management of any other services, including managed or specialized services,⁴ offered by the provider even if such services are provided over common facilities.** As demonstrated earlier, modern communications networks are used to provide a wide (and constantly changing) array of services, each of which is comprised of features and functionality sought by users and each of which accordingly places different demands on the network. In determining whether and how to provide these services, network operators weigh numerous factors and make complex calculations about whether to invest substantial amounts of capital to deploy additional capacity or manage within current capacity. Further, a balance made by one entity likely is irrelevant (or at best of minimal relevance) to a balance made by another – after all, today’s communications networks vary tremendously, even among providers offering the same or similar arrays of services.

- **Broadband Internet access providers may offer users different classes of services with varying performance capabilities, which may include the imposition of limitations on users’ bandwidth and consumption subject to full and understandable disclosure.** It is a common and well-accepted practice today for broadband providers to offer classes (tiers) of service, which are sold at different prices for different performance capabilities (e.g. bandwidth) and with different terms and conditions (e.g. limitations on consumption). The key is to make sure users are given

⁴ Managed or specialized services are those services provided by communications network operators where the operator takes control of the management and operation of the service from the user according to specific parameters and requirements agreed upon with the user and may include planning, optimization, and development functions, integration with the business services of the user, and support and maintenance activities.

sufficiently clear information about these classes of service – both capabilities and limitations – so they can make informed decisions.

- **Broadband Internet access providers may provide quality of service or other performance guarantees to providers and users of applications so long as they are not offered on an exclusive basis for any particular application category.** As discussed throughout this document, many Internet applications are highly-sensitive to latency and jitter and packet loss – and newer real-time video applications also must have sufficient bandwidth to provide a satisfactory experience for users. Providers of applications and users of these applications would benefit from QoS guarantees and, in fact, often request them. (Many of these providers already use CDNs to ensure QoS from the host server to a server collocated with the broadband Internet access provider.) So long as any QoS guarantees are offered on a non-exclusive basis – that is to all providers of any particular category of application – they should be permitted.

- **Broadband Internet access providers may choose to increase or alter network capacity by building additional facilities or installing or tuning electronics -- and the government can provide incentives to build additional capacity⁵ -- but providers should not be required to do so.** Decisions regarding increasing network capacity are exceedingly complex, involving not just network engineering concerns but all the factors that are part of a business case. In addition, these decisions often have to be made relatively rapidly given the growth of Internet traffic. Further, the circumstances that lead to the possible investment in more capacity may be driven by a small subset of subscribers. Finally, it would be unprecedented for the government to mandate or prevent the construction of facilities and equipment by communications providers other than those who have a monopoly.

- **Broadband Internet access providers may throttle and shape Internet access traffic (1) so long as they do so for all applications of a particular type and so long as the action is of a short duration and not repeatedly imposed, (2) if it is imposed on an individual user, so long as such action is part of the terms and conditions agreed upon by the user when choosing the service, or (3) in emergency situations .** Because of the potential high volatility of Internet traffic, network managers often face surges beyond those anticipated by and engineered into the network. To ensure a satisfactory level of service, network managers should be able to address these events in several ways. First, they may throttle or shape traffic if it applies to all applications of a particular type. Second, they may throttle or shape traffic for a particular user if the provider informs the user in clear and understandable terms in advance that such activity may occur. Third, they may throttle or shape traffic if

⁵ The FTTH Council, in fact, believes the government should be involved in providing incentives to accelerate deployment of FTTH and other next-generation access networks, and it has consistently advocated for measures to accomplish this objective, including investment tax credits and tax credit bonds.

emergency situations – to permit critical traffic to get through or to prevent the network from crashing.

- **Broadband Internet access providers may permit CDNs or others, including themselves or affiliated entities, to install and have traffic directed towards local caching equipment.** CDNs are so widespread and have such generally recognized benefits for users and content and applications providers that they should be encouraged – even though they may give advantages to particular entities (as opposed to particular applications).

- **Broadband Internet access providers may block or otherwise restrict spam, malware, and similar traffic harmful to the network and unlawful traffic (or unlawful transfer of content) – as well as provide priorities for emergency traffic and secure transmissions.** There is a virtually unanimous agreement that these activities are either so harmful in the case of spam and malware or beneficial in the case of emergency and secure traffic that network managers should be able to address them expeditiously by use of practices generally accepted in the industry.

- **Broadband Internet access providers should publicly disclose information about network usage limitations or conditions and general information about their network management practices but may limit disclosure of network management practices that would negate or otherwise undermine any permissible network management practice.** Providing users and applications and content providers with sufficient and understandable information about network usage limitations or conditions is important to ensure they make informed choices about subscribing to the service; however, providers should not be ordered to disclose information regarding the precise nature of any network management practices, including technical methods employed to manage traffic. Such granular detail about network management practices is often proprietary, crucial to preserve network security and reliability, and is constantly changing in response to new congestion problems or malware attacks.